

A Framework for Parallel and Distributed Training of Neural Networks

Simone Scardapane^{a,*}, Paolo Di Lorenzo^b

^a*Department of Information Engineering, Electronics and Telecommunications,
“Sapienza” University of Rome, Via Eudossiana 18, 00184 Rome, Italy*

^b*Department of Engineering, University of Perugia, Via G. Duranti 93, 06125, Perugia,
Italy*

Abstract

The aim of this paper is to develop a general framework for training neural networks (NNs) in a distributed environment, where training data is partitioned over a set of agents that communicate with each other through a sparse, possibly time-varying, connectivity pattern. In such distributed scenario, the training problem can be formulated as the (regularized) optimization of a non-convex social cost function, given by the sum of local (non-convex) costs, where each agent contributes with a single error term defined with respect to its local dataset. To devise a flexible and efficient solution, we customize a recently proposed framework for non-convex optimization over networks, which hinges on a (primal) convexification-decomposition technique to handle non-convexity, and a dynamic consensus procedure to diffuse information among the agents. Several typical choices for the training criterion (e.g., squared loss, cross entropy, etc.) and regularization (e.g., ℓ_2 norm, sparsity inducing penalties, etc.) are included in the framework and explored along the paper. Convergence to a stationary solution of the social non-convex problem is guaranteed under mild assumptions. Additionally, we show a principled way allowing each agent to exploit a multi-core architecture (e.g., a local cloud) in order to parallelize its local optimization step, resulting in strategies that are both distributed (across the agents) and parallel (inside each agent) in nature. A comprehensive set of experimental

*Corresponding author. Phone: +39 06 44585495, Fax: +39 06 4873300.

Email addresses: simone.scardapane@uniroma1.it (Simone Scardapane),
paolo.dilorenzo@unipg.it (Paolo Di Lorenzo)

results validate the proposed approach.

Keywords: Neural network, distributed learning, parallel computing, networks.

1. Introduction

We consider the problem of training a Neural Network (NN) model, when training data is distributed over different agents that are connected by a sparse, possibly time-varying, communication network. To grasp the main motivation, let us consider a ‘smart’ environment, wherein thousands of low-power sensors (e.g. cameras, wearables, etc.) are embedded to provide context-aware assistance, security provisioning, and so forth (Pottie and Kaiser, 2000; Boric-Lubeke and Lubecke, 2002). If the amount of produced data is small and we can count on a very reliable communication network, we may think of a centralized approach where all the data are transmitted to one (or more) fusion center that performs the learning task. However, in big data applications, sharing local information with a central processor might be either unfeasible or not economical/efficient, owing to the large size of the network and volume of data, time-varying network topology, energy constraints, robustness and/or privacy concerns. Performing the computation in a centralized fashion may raise robustness concerns as well, since the central processor represents a bottleneck and an isolated point of failure. For these reasons, effective learning methods must necessarily exploit distributed computation/learning architectures (with possibly parallelized multi-core processors), while keeping into account the distributed large-scale storage of data over the network and communication constraints. Very often, the implementation of such learning schemes requires the training of a *shared* predictive function. Considering the previous example, suppose that a set of embedded cameras is taking multiple high-resolution photos of a possible security threat. In this case, if the threat needs to be recognized quickly in the near future, the sensors need to train a common (shared) classifier that must leverage on all the currently acquired photos, in order to obtain a sufficiently high accuracy. These problems are ubiquitous in the real world, and appear in many practical systems such as, e.g., wireless sensor networks (Predd et al., 2006), smart grids, distributed databases (Lazarevic and Obradovic, 2002), robotic swarms, just to name a few.

If a predictive behavior is needed, however, the designer of the distributed

system has to answer a necessary question: What kind of model should be chosen as a classifier/regressor? Since deep NNs are currently obtaining state-of-the-art results in several fields (Schmidhuber, 2015; LeCun et al., 2015), employing them appears as a reasonable choice. Nevertheless, somewhat surprisingly, the literature on distributed training algorithms for NNs satisfying all the above requirements is extremely scarce. Most authors resort either to an ensemble of models trained independently by the agents (Lazarevic and Obradovic, 2002; Zhang and Zhong, 2013), or to strategies requiring the sum of the gradients’ contributions for all agents at every single iteration (Samet and Miri, 2012; Georgopoulos and Hasler, 2014), exploiting the additivity of the gradients updates. Both these approaches can be easily shown to be unsatisfactory in general. In the former case, we have no guarantee that the ensemble of models will perform as good as a single model trained on the collection of all local datasets. In the latter case, instead, a global sum at every iteration might be infeasible due to an excessive amount of communication, particularly for large models comprising several hundred thousands parameters. It is also worth mentioning that a lot of research has been devoted recently to the design of parallel, asynchronous versions of stochastic gradient descent for training NNs on large clusters of commodity hardware (Dean et al., 2012; Sak et al., 2014; Abadi et al., 2016). However, all these previous methods require the presence of at least one central server node, which coordinates the learning process; thus, they are not applicable in our context.

One of the reasons for the lack of distributed training methods for NNs is that, in principle, these methods require the solution of a distributed *non-convex* optimization problem, which was tackled only in a few papers even in the optimization literature (Bianchi and Jakubowicz, 2013; Di Lorenzo and Scutari, 2016). On the other side, if we turn our attention to methods for convex learning problems, the literature on their distributed training is vast, including algorithms for decentralized optimization of linear predictors (Xiao et al., 2007; Sayed, 2014; Sayed et al., 2014), sparse linear models (Mateos et al., 2010; Di Lorenzo and Sayed, 2013), kernel ridge regression (Predd et al., 2006, 2009), random-weights networks (Huang and Li, 2015; Scardapane et al., 2015, 2016c), support vector machines (Navia-Vázquez et al., 2006; Lu et al., 2008; Forero et al., 2010; Scardapane et al., 2016b), and kernel filtering (Perez-Cruz and Kulkarni, 2010; Gao et al., 2015).

Contribution: In this paper, we propose an algorithmic framework for training general NN models in a fully distributed scenario, which encom-

passes several common loss functions and regularization terms.¹ In particular, we build upon the recently proposed in-network nonconvex optimization (NEXT) algorithm (Di Lorenzo and Scutari, 2016), i.e., one of the first methods to solve distributed non-convex optimization problems over networks of agents. The algorithm, which leverages on the so-called successive convex approximation (SCA) family of methods (Facchinei et al., 2015), is built upon two foundational ideas. First and foremost, at every iteration, the original non-convex problem is replaced with a strongly convex approximation, which is solved *locally* at every agent. As we will illustrate along the paper, several kinds of convexification are possible, resulting in different trade-offs in terms of computational complexity and speed of convergence. Second, the framework exploits a dynamic consensus procedure (Zhu and Martínez, 2010), so that each agent can recover the information relative to all the other agents, which typically is not available at its local side. The resulting algorithms are shown to be convergent to a stationary solution of the social non-convex problem under loose requirements relative to the agents’ communication topology, the choice of the algorithm’s parameters, and the structure of the optimization problem. A further interesting aspect of the framework presented here is that the local optimization problems can be easily parallelized in a principled way (up to one NN parameter per available processor), without losing the convergence properties of the framework. Consider, for example, the case of multiple medical institutions requiring the training of a common NN (e.g., for diagnosis purposes) leveraging on all historical clinical information (Vieira-Marques et al., 2006). In this case, a decentralized algorithm is required due to strong privacy concerns on the release of medical, sensible information about the patients. Nonetheless, each institution may have access to an internal private cloud infrastructure. Using the framework outlined in this paper, privacy is guaranteed via the use of a distributed protocol, while each institution can parallelize its optimization steps using local cloud computing hardware. In this way, the resulting algorithms are both distributed (across the nodes) and parallel (inside each node) in nature. A comprehensive set of experimental results validate the proposed approach.

Outline of the paper: The rest of the paper is organized as follows. In Section 2, we formalize the problem of distributed NN training. In Section

¹A preliminary version of this work, focusing only on the squared loss function, was presented in (Di Lorenzo and Scardapane, 2016).

3 we describe the general framework for distributed NN training built upon the NEXT algorithm. Then, in Section 4, we consider the customization of the framework to different loss functions (squared loss, cross entropy, etc.) and regularization terms (ℓ_2 norm, sparsity inducing penalties, etc.). Section 5 describes a principled way to parallelize the optimization phase. In Section 6.1, we perform a large set of experiments aimed at assessing the performance of the proposed framework. Finally, Section 7 draws some conclusions and future lines of research.

Notation: We denote vectors using boldface lowercase letters, e.g., \mathbf{a} ; matrices are denoted by boldface uppercase letters, e.g., \mathbf{A} . All vectors are assumed to be column vectors. The operator $\|\cdot\|_p$ is the standard ℓ_p norm on an Euclidean space. For $p = 2$, it coincides with the Euclidean norm, while for $p = 1$ we obtain the Manhattan (or taxicab) norm defined for a generic vector $\mathbf{v} \in \mathbb{R}^B$ as $\|\mathbf{v}\|_1 = \sum_{k=1}^B |v_k|$. The notation $a[n]$ denotes the dependence of a on the time-index n . Other notation is introduced along the paper when required.

2. Problem Formulation

Let us consider the problem of training a generic NN model $f(\mathbf{w}; \mathbf{x})$, where $\mathbf{x} \in \mathbb{R}^d$ denotes the d -dimensional input vector of the network, whereas $\mathbf{w} \in \mathbb{R}^Q$ is the vector collecting all the adaptable parameters that we aim to optimize. Note that we are considering the NN as a function of its parameters, as this will make the following derivation simpler. We are not concerned with the specific structure of the NN $f(\cdot)$ (i.e., number of hidden layers, choice of the activation functions, etc.), as long as the following assumptions are satisfied for any possible input vector $\mathbf{x} \in \mathbb{R}^d$.

Assumption A [On the NN model]:

- (A1) f is in C^1 , i.e., it is continuously differentiable with respect to \mathbf{w} ;
- (A2) f has Lipschitz continuous gradient, with respect to \mathbf{w} , for some Lipschitz constant L , i.e.:

$$\|\nabla_{\mathbf{w}} f(\mathbf{w}_1; \mathbf{x}) - \nabla_{\mathbf{w}} f(\mathbf{w}_2; \mathbf{x})\|_2 \leq L \|\mathbf{w}_1 - \mathbf{w}_2\|_2. \quad (1)$$

Assumption A is satisfied by most NN models commonly used in the literature, with the only notable exception of NN having non-differentiable

activation functions such as ReLu neurons (Glorot et al., 2011), maxout neurons (Goodfellow et al., 2013), and a few others. Nonetheless, convergence guarantees for these architectures are relatively uncommon even in the centralized case. In this paper, we are concerned with distributed architectures, where the data required to train the NN is not available on a centralized location, but is instead partitioned among I interconnected agents. Prototypical examples of agents can be sensors in a wireless sensor networks (WSN), peers in a P2P network, power units in a smart grid, or mobile robots in a robotic swarm. At every specific time instant n , the communication network enabling interaction among the agents is modeled as a directed graph (digraph) $\mathcal{G}[n] = (\mathcal{V}, \mathcal{E}[n])$, where $\mathcal{V} = \{1, \dots, I\}$ is the vertex set (i.e., the set of agents), and $\mathcal{E}[n]$ is the set of (possibly) time-varying directed edges. The in-neighborhood of agent i at time n (including node i) is defined as $\mathcal{N}_i^{\text{in}}[n] = \{j | (j, i) \in \mathcal{E}[n]\} \cup \{i\}$: node i can receive information from node $j \neq i$ at time instant n only if $j \in \mathcal{N}_i^{\text{in}}[n]$. Due to this, each agent has a limited view and knowledge about the overall (possibly time-varying) network. Also, we assume that there is no agent (or finite number of them) that is able to collect all the data and coordinate the overall learning process. Associated with each graph $\mathcal{G}[n]$, we introduce (possibly) time-varying weights $c_{ij}[n]$ matching $\mathcal{G}[n]$:

$$c_{ij}[n] = \begin{cases} \theta_{ij} \in [\vartheta, 1] & \text{if } j \in \mathcal{N}_i^{\text{in}}[n]; \\ 0 & \text{otherwise,} \end{cases} \quad (2)$$

for some $\vartheta \in (0, 1)$, and define the matrix $\mathbf{C}[n] \triangleq (c_{ij}[n])_{i,j=1}^I$. These weights will be used later on in the definition of the proposed algorithm, in order to locally combine the information diffused over every neighborhood. For the purpose of training the NN, we assume that the i th agent has access to a local training dataset of N_i examples, denoted as $\mathcal{S}_i = \{\mathbf{x}_{i,m}, d_{i,m}\}_{m=1}^{N_i}$, where we consider a single-output problem with $d_{i,m} \in \mathbb{R}$ for simplicity of overall notation. The output of the NN is an integer or a real value, depending on whether we are facing a classification task or a regression task, respectively. Given all the previous definitions, a general formulation for the distributed training of NNs can be cast as the minimization of a social cost function G

plus a regularization term $r(\cdot)$, which writes as:

$$\min_{\mathbf{w}} U(\mathbf{w}) = G(\mathbf{w}) + r(\mathbf{w}) = \sum_{i=1}^I g_i(\mathbf{w}) + r(\mathbf{w}), \quad (3)$$

where $g_i(\cdot)$ is the error term relative to the i th local dataset:

$$g_i(\mathbf{w}) = \sum_{m \in \mathcal{S}_i} l(d_{i,m}, f(\mathbf{w}; \mathbf{x}_{i,m})), \quad (4)$$

with $l(\cdot, \cdot)$ denoting a generic (convex) loss function, while $r(\mathbf{w})$ is a regularization term. Due to the nonlinearity of the NN model $f(\mathbf{w}; \mathbf{x})$, problem (3) is typically *non-convex*. In this work, we consider the following assumptions on the functions involved in (3)-(4).

Assumption B [On Problem (3)]:

- (B1) l is convex and C^1 , with Lipschitz continuous gradient;
- (B2) r satisfies (B1), or it is a nondifferentiable convex function with bounded subgradients;
- (B3) U is coercive, i.e., $\lim_{\|\mathbf{w}\| \rightarrow \infty} U(\mathbf{w}) = +\infty$.

The structure of the function l in (4) depends on the learning task (i.e., regression, classification, etc.). Typical choices are the squared loss for regression problems, and the cross-entropy for classification tasks (Haykin, 2009). The regularization function $r(\mathbf{w})$ in (3) is commonly chosen to avoid overfitted solutions and/or impose a specific structure in the solution, e.g. sparsity or group sparsity. Typical choices are the ℓ_2 and ℓ_1 norms. All these functions satisfy Assumption B, and will be discussed in detail in the sequel. In view of the distributed nature of the problem, the i th agent knows its own cost function g_i and the common regularization term r , but it does not have access to g_j for $j \neq i$, nor can it exchange freely its own dataset \mathcal{S}_i due to a variety of reasons, including privacy, data volume, and communication constraints. This aspect, combined with the non-convexity of (3), makes optimizing (3) in a distributed fashion a challenging problem, which has no ready-to-use solution available in the literature. The design of such algorithmic framework is the topic of the next three sections.

3. NEXT: In-Network Successive Convex Approximation

In this section, we review the basics of the NEXT framework proposed in (Di Lorenzo and Scutari, 2016), which was designed to solve general nonconvex distributed problems of the form (3). The next section will then focus on how to customize the framework to the NN distributed training problem considered in this paper. Due to lack of space, we provide only a very brief introduction to the NEXT framework, and we refer the interested reader to (Di Lorenzo and Scutari, 2016) for a full treatment, which also includes a proof of the convergence results. NEXT combines SCA techniques (Step 1) with dynamic consensus mechanisms (Steps 2 and 3), as described next.

Step 1 (local SCA optimization): Each agent i maintains a local estimate $\mathbf{w}_i[n]$ of the optimization variable \mathbf{w} that is iteratively updated. Solving directly Problem (3) may be too costly (due to the nonconvexity of G) and is not even feasible in a distributed setting. One may then prefer to approximate Problem (3), in some suitable sense, in order to permit each agent to compute *locally* and *efficiently* the new iteration. In particular, writing $G(\mathbf{w}_i) = g_i(\mathbf{w}_i) + \sum_{j \neq i} g_j(\mathbf{w}_i)$, we consider a convexification of G having the following form: i) at every iteration n , the (possibly) nonconvex $g_i(\mathbf{w}_i)$ is replaced by a strongly convex surrogate, say $\tilde{g}_i(\cdot; \mathbf{w}_i[n]) : \mathbb{R}^Q \rightarrow \mathbb{R}$, which may depend on the current iterate $\mathbf{w}_i[n]$; and ii) $\sum_{j \neq i} g_j(\mathbf{w}_i)$ is linearized around $\mathbf{w}_i[n]$. More formally, the proposed updating scheme reads: at every iteration n , given the local estimate $\mathbf{w}_i[n]$, each agent i solves the *strongly convex* optimization problem:

$$\begin{aligned} \tilde{\mathbf{w}}_i[n] &= \arg \min_{\mathbf{w}_i} \tilde{U}_i(\mathbf{w}_i; \mathbf{w}_i[n], \boldsymbol{\pi}_i[n]) \\ &= \arg \min_{\mathbf{w}_i} \tilde{g}_i(\mathbf{w}_i; \mathbf{w}_i[n]) + \boldsymbol{\pi}_i[n]^T (\mathbf{w}_i - \mathbf{w}_i[n]) + r(\mathbf{w}_i). \end{aligned} \quad (5)$$

where

$$\boldsymbol{\pi}_i[n] \triangleq \sum_{j \neq i} \nabla_{\mathbf{w}} g_j(\mathbf{w}_i[n]). \quad (6)$$

The evaluation of (6) would require the knowledge of all $\nabla g_j(\mathbf{w}_i[n])$, $j \neq i$ at node i . This information is not directly available at node i ; we will cope with this local lack of global knowledge later on in step 3. Once the surrogate problem (5) is solved, each agent computes an auxiliary variable, say $\mathbf{z}_i[n]$,

as the convex combination:

$$\mathbf{z}_i[n] = \mathbf{w}_i[n] + \alpha[n] (\tilde{\mathbf{w}}_i[n] - \mathbf{w}_i[n]) , \quad (7)$$

where $\alpha[n]$ is a possibly time-varying step-size sequence. This concludes the optimization phase of the algorithm. An appropriate choice of the surrogate function $\tilde{g}_i(\cdot; \mathbf{w}_i[n])$ guarantees the coincidence between the fixed-points of $\tilde{\mathbf{w}}_i[n]$ and the stationary solutions of Problem (3). The main results are given in the following proposition:

Proposition 1. *Given Problem (3) under A1-A2 and B1-B3, suppose that \tilde{g}_i satisfies the following conditions:*

(F1) $\tilde{g}_i(\cdot; \mathbf{w})$ is uniformly strongly convex with $\tau_i > 0$;

(F2) $\nabla \tilde{g}_i(\mathbf{w}; \mathbf{w}) = \nabla g_i(\mathbf{w})$ for all \mathbf{w} ;

(F3) $\nabla \tilde{g}_i(\mathbf{w}; \cdot)$ is uniformly Lipschitz continuous.

Then, the set of fixed-point of $\tilde{\mathbf{w}}_i[n]$ in (5) coincides with that of the stationary solutions of (3).

Conditions F1-F3 state that \tilde{g}_i should be regarded as a convex approximation of g_i at the point \mathbf{w} , which preserves the first order properties of g_i . Several feasible choices are possible for a given g_i ; the appropriate one depends on computational and communication requirements. The goal of the next section will be to illustrate some possible choices for the local surrogate cost \tilde{g}_i properly customized to our distributed NN training problem.

Step 2 (agreement update): To force the asymptotic agreement among the \mathbf{w}_i 's, a consensus-based step is employed on the auxiliary variables $\mathbf{z}_i[n]$'s. Each agent i updates its local variable $\mathbf{w}_i[n]$ as:

$$\mathbf{w}_i[n+1] = \sum_{j \in \mathcal{N}_i^{\text{in}}[n]} c_{ij}[n] \mathbf{z}_j[n], \quad (8)$$

where $\mathbf{C}[n] = (c_{ij}[n])_{ij}$ is defined in (2), and satisfies

$$\mathbf{C}[n] \mathbf{1} = \mathbf{1} \quad \text{and} \quad \mathbf{1}^T \mathbf{C}[n] = \mathbf{1}^T \quad \forall n. \quad (9)$$

Since the weights are constrained by the network topology, (8) can be implemented via local message exchanges: agent i updates its estimate \mathbf{w}_i by

averaging over the current solutions $\mathbf{z}_j[n]$ received from its neighbors. The double stochasticity condition in (9) can be achieved according to a variety of predefined strategies, including the Metropolis-Hastings criterion (Xiao et al., 2007), or by optimizing a cost function with respect to the spectral properties of the graph (Xiao and Boyd, 2004).

Step 3 (diffusion of information over the network): The computation of $\tilde{\mathbf{w}}_i[n]$ in (5) is not fully distributed yet, because the evaluation of $\boldsymbol{\pi}_i[n]$ in (6) would require the knowledge of all $\nabla g_j(\mathbf{w}_i[n])$, $j \neq i$, which is a global information that is not available locally at node i . To cope with this issue, as proposed in (Di Lorenzo and Scutari, 2016), we replace $\boldsymbol{\pi}_i[n]$ in (5) with a *local* estimate, say $\tilde{\boldsymbol{\pi}}_i[n]$, asymptotically converging to $\boldsymbol{\pi}_i[n]$. Thus, we can update the local estimate $\tilde{\boldsymbol{\pi}}_i[n]$ in a fully distributed manner as:

$$\tilde{\boldsymbol{\pi}}_i[n] \triangleq I \cdot \mathbf{y}_i[n] - \nabla g_i(\mathbf{w}_i[n]), \quad (10)$$

where $\mathbf{y}_i[n]$ is a local auxiliary variable (controlled by user i) that aims to asymptotically track the average of the gradients. This can be done updating $\mathbf{y}_i[n]$ according to the following dynamic consensus recursion:

$$\mathbf{y}_i[n+1] \triangleq \sum_{j=1}^I c_{ij}[n] \mathbf{y}_j[n] + (\nabla g_i(\mathbf{w}_i[n+1]) - \nabla g_i(\mathbf{w}_i[n])) \quad (11)$$

with $\mathbf{y}_i[0] \triangleq \nabla_{\mathbf{w}_i} g_i(\mathbf{w}_i[0])$. Note that the update of $\mathbf{y}_i[n]$ and thus $\tilde{\boldsymbol{\pi}}_i[n]$ can be now performed locally with message exchanges with the agents in the neighborhood.

The overall procedure is summarized in Algorithm 1, where $\nabla \mathbf{g}_i[n]$ is used as a simplified notation for $\nabla_{\mathbf{w}_i} g_i(\mathbf{w}_i[n])$. Its convergence properties are reported in the following Theorem.

Theorem 2. *Let $\{\mathbf{w}[n]\}_n \triangleq \{(\mathbf{w}_i[n])_{i=1}^I\}_n$ be the sequence generated by Algorithm 1, and let $\{\bar{\mathbf{w}}[n]\}_n \triangleq \{(1/I) \sum_{i=1}^I \mathbf{w}_i[n]\}_n$ be its average. Suppose that i) Assumptions A and B hold; ii) the sequence of graphs describing the network is B -strongly connected²; iii) condition (9) holds; and iii) the step-size*

²Formally, there exists an integer $B > 0$ such that the graph $\mathcal{G}[k] = (\mathcal{V}, \mathcal{E}_B[k])$, with $\mathcal{E}_B[k] = \bigcup_{n=kB}^{(k+1)B-1} \mathcal{E}[n]$ is strongly connected, for all $k \geq 0$.

sequence $\{\alpha[n]\}_n$ is chosen so that $\alpha[n] \in (0, 1]$ for all n and:

$$\sum_{n=0}^{\infty} \alpha[n] = \infty \quad \text{and} \quad \sum_{n=0}^{\infty} \alpha[n]^2 < \infty. \quad (12)$$

Then, if the sequence $\{\bar{\mathbf{w}}[n]\}_n$ is bounded, (a) all its limit points are stationary solutions of (3); (b) all the sequences $\{\mathbf{w}_i[n]\}_n$ asymptotically agree, i.e., $\|\mathbf{w}_i[n] - \bar{\mathbf{w}}[n]\|_2 \xrightarrow[n \rightarrow \infty]{} 0$, for all i .

Proof. Under the above assumptions on the NN model in (3), the network among agents, the algorithm's parameters, and the boundness of $\{\bar{\mathbf{w}}[n]\}_n$, all conditions of Theorem 3 in (Di Lorenzo and Scutari, 2016) are satisfied, and the convergence result follows. \square

It is interesting to notice that convergence conditions are particularly loose. With respect to the network connecting the agents, it is enough to ensure connectivity over a finite (but arbitrary) *union* of time instants. Step-size sequences satisfying (12) can be derived easily, e.g., the following quadratically decreasing rule was found to be particularly effective in our experiments:

$$\alpha[n] = \alpha[n-1] (1 - \varepsilon \alpha[n-1]) , \quad (14)$$

where $\alpha[0], \varepsilon \in (0, 1]$ must be chosen by the user. Boundedness of the sequence $\{\bar{\mathbf{w}}[n]\}_n$ is not a concern in general, since it can always be achieved by imposing a finite (but arbitrarily large) box constraint during the execution of the algorithm. Typically, if the box is chosen sufficiently large, its effect can be neglected in practical cases, because the algorithm will end up converging to a local solution of (3) independently on the effect of the box.

The per-iteration cost of the algorithm is clearly dominated by the solution of the surrogate optimization problem in (13). As we will see in the next section, the flexibility of the framework allows to select different choices of surrogate functions, typically impacting the complexity/performance trade-off of the algorithm. The framework can be accelerated in two ways. First, we can parallelize the surrogate optimization in (13); this point will be discussed in Section 5. Second, at each iteration n , we can consider an *inexact* solution of the surrogate problems in (13) within a user-specified error bound $\epsilon_i[n]$. In this case, it can be shown that convergence is still guaranteed, as long as the following condition is satisfied: $\sum_{n=0}^{\infty} \alpha[n] \epsilon_i[n] < \infty$, $\forall i \in 1, \dots, I$, which establish a decaying rate with which the error sequence must diminish over

Algorithm 1: NEXT Framework for Distributed Optimization of (3)

Data : $\mathbf{w}_i[0], \mathbf{y}_i[0] = \nabla g_i[0], \boldsymbol{\pi}_i[0] = I\mathbf{y}_i[0] - \nabla g_i[0], \forall i = 1, \dots, I$, and $\{\mathbf{C}[n]\}_n$. Set $n = 0$.

(S.1) If $\mathbf{w}_i[n]$ satisfies a global termination criterion: STOP;

(S.2) **Local Optimization:** Each agent

(a) computes $\tilde{\mathbf{w}}_i[n]$ as:

$$\tilde{\mathbf{w}}_i[n] = \arg \min_{\mathbf{w}_i} \tilde{U}_i(\mathbf{w}_i; \mathbf{w}_i[n], \tilde{\boldsymbol{\pi}}_i[n]), \quad (13)$$

(b) updates its local variable $\mathbf{z}_i[n]$:

$$\mathbf{z}_i[n] = \mathbf{w}_i[n] + \alpha[n] (\tilde{\mathbf{w}}_i[n] - \mathbf{w}_i[n]).$$

(S.3) **Consensus update:** Each agent i collects data from its current neighbors and updates the estimates as:

$$\begin{aligned} \text{(a)} \quad \mathbf{w}_i[n+1] &= \sum_{j=1}^I c_{ij}[n] \mathbf{z}_j[n], \\ \text{(b)} \quad \mathbf{y}_i[n+1] &= \sum_{j=1}^I c_{ij}[n] \mathbf{y}_j[n] + (\nabla g_i[n+1] - \nabla g_i[n]), \\ \text{(c)} \quad \tilde{\boldsymbol{\pi}}_i[n+1] &= I \cdot \mathbf{y}_i[n+1] - \nabla g_i[n+1]. \end{aligned}$$

(S.4) $n \leftarrow n + 1$, and go to (S.1).

time. For further details, we refer to (Di Lorenzo and Scutari, 2016, Theorem 4).

4. Strategies for Distributed NN Training

In this section, we customize the NEXT framework for the solution of several distributed NN training problems. In particular, we focus on the choice of the surrogate functions \tilde{g}_i in (5). In particular, from Proposition 1, we know that they must be chosen to satisfy F1-F3. Thus, we explore two general-purpose strategies that can be used to this end, before analyzing

some practical algorithms resulting from the combination of these two strategies with common choices of the loss function and the regularization term. Essentially, the aim of $\tilde{g}_i(\cdot)$ is to provide a strongly convex approximation of (the non-convex) g_i around the current point, preserving (at least) the first-order information of the original function. Then, the most basic idea is to linearize the entire g_i , irrespectively of the actual choice of loss function l , as:

$$\begin{aligned}\tilde{g}_i^{\text{FL}}(\mathbf{w}_i; \mathbf{w}_i[n]) &= g_i(\mathbf{w}_i[n]) + \nabla g_i(\mathbf{w}_i[n])^T (\mathbf{w}_i - \mathbf{w}_i[n]) \\ &\quad + \frac{\tau}{2} \|\mathbf{w}_i - \mathbf{w}_i[n]\|_2^2, \end{aligned} \quad (15)$$

where the last term in (15) is a proximal regularization term (with $\tau \geq 0$) used to ensure strong convexity; in what follows, we will refer to (15) as the full linearization strategy (FL). In general, the use of the FL strategy leads to the formulation of surrogate problems in (13) allowing for a simple, closed-form solution for most choices of regularization. At the same time, this strategy is throwing away most information of $g_i(\cdot)$, by only keeping first-order information on its gradient. For this reason, the resulting family of algorithms can possess a slow convergence speed, similarly to what happens with the use of (centralized) steepest descent optimization procedures.

To implement a more sophisticated approximation aimed at preserving the hidden convexity in the problem, we start noticing that the loss function in (4) is composed of the summation of terms, each one given by the composition of an exterior convex function (i.e., the loss function l), and an interior nonlinear function (i.e., the NN model f). Then, a possible choice for \tilde{g}_i is to preserve the convexity of l , while linearizing f around the current estimate $\mathbf{w}_i[n]$, and a generic input point $\mathbf{x}_{i,m}$, as:

$$\tilde{f}(\mathbf{w}_i; \mathbf{w}_i[n], \mathbf{x}_{i,m}) = f(\mathbf{w}_i[n]; \mathbf{x}_{i,m}) + \nabla f(\mathbf{w}_i[n]; \mathbf{x}_{i,m})^T (\mathbf{w}_i - \mathbf{w}_i[n]). \quad (16)$$

Then, the surrogate \tilde{g}_i is obtained as:

$$\tilde{g}_i^{\text{PL}}(\mathbf{w}_i; \mathbf{w}_i[n]) = \sum_{m \in \mathcal{S}_i} l(d_{i,m}, \tilde{f}(\mathbf{w}_i; \mathbf{w}_i[n], \mathbf{x}_{i,m})) + \frac{\tau}{2} \|\mathbf{w}_i - \mathbf{w}_i[n]\|^2, \quad (17)$$

with $\tau \geq 0$. We will refer to (17) as the partial linearization (PL) strategy. It is straightforward to check that the surrogate \tilde{g}_i^{PL} in (17) satisfies the properties F1-F3 required by Proposition 1.

In the remainder of the section, we consider a set of practical examples resulting from the use of our general framework.

4.1. Case 1: ridge regression cost

As a first example, we consider the use of a squared loss function combined with a classical ℓ_2 norm regularization on the weights (also known as weight decay in the NN literature (Moody et al., 1995)):

$$l(a, b) \triangleq (a - b)^2, \quad r(\mathbf{w}) \triangleq \frac{\lambda}{2} \|\mathbf{w}\|_2^2, \quad (18)$$

where λ is a positive regularization parameter. Historically, this is the most common training criterion for NNs, and it is still widely used today for regression problems. Being equivalent to a nonlinear ridge regression, we borrow this terminology here.

Let us begin with the FL strategy in (15). Note that, thanks to the specific form of the regularizer, the resulting optimization problem in (13) is already strongly convex, so that we can set $\tau = 0$. Then, using (15) and (18), the surrogate problem in (13) reduces to the minimization of a positive definite quadratic function, which admits a simple closed form solution, given by:

$$\boxed{\tilde{\mathbf{w}}_i[n] = -\frac{1}{\lambda} \left(\nabla \mathbf{g}_i[n] + \tilde{\boldsymbol{\pi}}_i[n] \right)}, \quad (19)$$

where as before $\nabla \mathbf{g}_i[n]$ is used as a simplified notation for $\nabla_{\mathbf{w}_i} g_i(\mathbf{w}_i[n])$. Eq. (19) represents the first practical implementation of the framework in Algorithm 1 for distributed NN training. As we can notice from (19), the FL strategy discards all information on the global cost function U in (3), except for a first-order approximation. Thus, the descent direction in (19) will be proportional to the opposite of the gradient of U , thanks to the current estimate $\tilde{\boldsymbol{\pi}}_i[n]$ of (6) that is locally available at node i . As we will see in the numerical results, the performance of the resulting distributed scheme is similar to a centralized gradient method, sharing its advantages (low computational complexity) and its drawbacks (possible slow convergence speed).

We now proceed considering the PL strategy in (17). To this aim, let us introduce the following ‘residual’ terms:

$$r_{i,m}[n] = y_{i,m} - f(\mathbf{w}_i[n]; \mathbf{x}_{i,m}) + \mathbf{J}_{i,m}[n]^T \mathbf{w}_i[n], \quad (20)$$

where $\mathbf{J}_{i,m}[n] = \nabla_{\mathbf{w}_i} f(\mathbf{w}_i[n]; \mathbf{x}_{i,m})$ is a Q -dimensional vector containing the derivatives of the NN output with respect to any single weight parameter. In the general case, it will be a matrix with one column per NN output. This quantity is sometimes denoted as the weight Jacobian (Blackwell, 2012), since it measures the influence of a small parameter change on the output of the neural network.³ Now, using (17) in (13), it is easy to show that the surrogate problem can be written again as the minimization of a positive definite quadratic form, given by:

$$\tilde{\mathbf{w}}_i[n] = \arg \min_{\mathbf{w}_i} \mathbf{w}_i^T \left(\mathbf{A}_i[n] + \frac{\lambda}{2} \mathbf{I} \right) \mathbf{w}_i - 2\mathbf{b}_i[n]^T \mathbf{w}_i, \quad (21)$$

where

$$\mathbf{A}_i[n] = \sum_{m \in S_i} \mathbf{J}_{i,m}[n] \mathbf{J}_{i,m}[n]^T, \quad (22)$$

$$\mathbf{b}_i[n] = \sum_{m \in S_i} \mathbf{J}_{i,m}[n] r_{i,m}[n] - 0.5 \tilde{\boldsymbol{\pi}}_i[n]. \quad (23)$$

As an interesting side note, in the NN literature the matrix (22) is known as an outer product approximation to the Hessian matrix of $g_i(\cdot)$ (i.e., the error function local to agent i), which is obtained by assuming that the error is uncorrelated with the second derivative of the network's output (Bishop, 2006, Section 5.4.2). Finally, solving the resulting minimization problem in (21), the solution $\tilde{\mathbf{w}}_i[n]$ of the surrogate problem, to be used in (13), is given by:

$$\tilde{\mathbf{w}}_i[n] = \left(\mathbf{A}_i[n] + \frac{\lambda}{2} \mathbf{I} \right)^{-1} \mathbf{b}_i[n]. \quad (24)$$

Differently from the FL strategy, whose computational complexity is linear in the number of parameters, in this case solving the surrogate problem is of the order $\mathcal{O}(Q^3)$, where Q is the number of adaptable NN parameters, due to the matrix inversion step. Nevertheless, as we will see in the numerical results, the resulting descent direction provides a very large improvement in terms of convergence speed. Additionally, this strategy can benefit from a larger relative speedup when employing the parallelization strategy described

³Note that a single back-propagation step per iteration is needed to build the weight Jacobian, as discussed in Bishop (2006, Section 5.3.4).

in Section 5.

4.2. Case 2a: squared error with weight sparsity

As a second example, let us consider again the use of a squared loss term l in (18), combined this time with a sparsity promoting term given by the ℓ_1 norm on the weight vector, i.e.,

$$r(\mathbf{w}) \triangleq \lambda \|\mathbf{w}\|_1 = \sum_{k=1}^Q |w_k|. \quad (25)$$

The ℓ_1 norm promotes sparsity of the weight vector, acting as a convex approximation of the non-convex, non-differentiable ℓ_0 norm (Tibshirani, 1996). While there exists customized algorithms to solve non-convex ℓ_1 regularized problems (Ochs et al., 2015), it is common in the NN literature to apply first-order procedures (e.g., stochastic descent with momentum) followed by a thresholding step to obtain sparse solutions (Bengio, 2012). In what follows, we illustrate the customization of the NEXT framework to this use case, using both FL and PL strategies in (15) and (17), respectively. In the FL case, using (15) and (25), with $\tau > 0$ to ensure strong convexity, the problem in (13) can be written as the minimization of the sum of q independent functions, as follows:

$$\begin{aligned} \tilde{\mathbf{w}}_i[n] = \arg \min_{\mathbf{w}_i} \sum_{k=1}^q & \left\{ (\nabla g_{ik}[n] + \tilde{\pi}_{ik}[n] - \tau w_{ik}[n]) w_{ik} \right. \\ & \left. + \frac{\tau}{2} w_{ik}^2 + \lambda |w_{ik}| \right\}. \end{aligned} \quad (26)$$

After some easy calculations, the solution of the optimization problem in (26) is given by the closed form expression:

$$\tilde{\mathbf{w}}_i[n] = \mathcal{S}_{\lambda/\tau} \left(\mathbf{w}_i[n] - \frac{1}{\tau} \nabla \mathbf{g}_i[n] - \frac{1}{\tau} \tilde{\boldsymbol{\pi}}_i[n] \right), \quad (27)$$

where

$$\mathcal{S}_\gamma(z) = \text{sign}(z) \max(0, |z| - \gamma), \quad (28)$$

is the (component-wise) soft thresholding function.

In the PL case, using (15) and (25), the problem in (13) can be cast as an ℓ_1 regularized quadratic program, given by:

$$\tilde{\mathbf{w}}_i[n] = \arg \min_{\mathbf{w}_i} \left\{ \mathbf{w}_i^T \left(\mathbf{A}_i[n] + \frac{\tau}{2} \mathbf{I} \right) \mathbf{w}_i - 2 (\mathbf{b}_i[n] + 0.5\tau \mathbf{w}_i[n])^T \mathbf{w}_i + \lambda \|\mathbf{w}_i\|_1 \right\}, \quad (29)$$

where $\mathbf{A}_i[n]$ and $\mathbf{b}_i[n]$ are given by (22) and (23), respectively. This is the first case we encounter where the solution of the optimization step cannot be expressed immediately in a closed form. Nevertheless, problem (29) is given by the sum of a strongly convex function and an ℓ_1 term, and many efficient strategies can be used for its approximate solution, including FISTA (Beck and Teboulle, 2009), coordinate descent methods (Cevher et al., 2014), and several others.

4.3. Case 2b: group sparse penalization

The formulation introduced in Sec. 4.2 can be easily extended to handle a *group sparse* penalization, which allows the selective removal of entire neurons during the training process, see, e.g., (Scardapane et al., 2016a). The basic idea is to force all the outgoing weights from a neuron to be *simultaneously* either non-zero or zero; the latter resulting in the direct removal of the neuron itself. Note that a neuron here can correspond to an input neuron, to a neuron in a hidden layer, or to a bias term, thus allowing the removal of input features, hidden neurons, and bias terms from the trained network (see (Scardapane et al., 2016a) for details). To this aim, let us suppose that the neurons are ordered and indexed as $1, \dots, P$. Also, let us denote by $\mathbf{w}_{i,p}$, $p = 1, \dots, P$, the subset of weights connecting the p th neuron with all the neurons in the following layer. Group sparsity can then be imposed by choosing in (3) the following regularization term:

$$r(\mathbf{w}) \triangleq \lambda \sum_{p=1}^P \rho_p \|\mathbf{w}_p\|_2, \quad (30)$$

where $\rho_p = \sqrt{\dim(\mathbf{w}_p)}$ are positive constants, $p = 1, \dots, P$, and $\dim(\mathbf{w}_p)$ denotes the dimension of vector space where \mathbf{w}_p lies.

Let us now analyze the customization of our framework when the FL strategy in (15) is applied. Then, let us define $\mathbf{a}_i[n] = \nabla \mathbf{g}_i[n] + \tilde{\boldsymbol{\pi}}_i[n] - \tau \mathbf{w}_i[n]$, denoting with $\mathbf{a}_{i,p}$ the restriction of \mathbf{a}_i to the indexes associated with the p th group. Thus, the surrogate problem in (13) writes as:

$$\tilde{\mathbf{w}}_i[n] = \arg \min_{\mathbf{w}_i} \sum_{p=1}^P \left\{ \mathbf{a}_{i,p}^T \mathbf{w}_{i,p} + \frac{\tau}{2} \|\mathbf{w}_{i,p}\|_2^2 + \lambda \rho_p \|\mathbf{w}_{i,p}\|_2 \right\}. \quad (31)$$

As we can notice from (31), the cost function is given by a summation of costs, each one dependent on a single neuron. Also in this case, even if problem (31) cannot be solved in closed form, it is possible to implement very fast and efficient algorithms for its solution, see, e.g., (Schmidt, 2010; Cevher et al., 2014). Furthermore, in the case each agent has a multi-core architecture, the structure of (31) makes straightforward the parallelization of computation, where each local processor can take care only of a subset of neurons. Finally, considering the PL strategy, the resulting formulation is equivalent to (29), with the only difference that (30) replaces the ℓ_1 norm. Again, many of the techniques mentioned before can be used to solve also the resulting (group sparse) strongly convex problem.

4.4. Case 3: cross-entropy loss

As an additional example, let us consider the case of binary classification, i.e., $d_{i,m} \in \{0, 1\}$. Then, assuming the output of the NN is limited between 0 and 1, a standard optimization criterion involves the cross-entropy loss function in (3), i.e.:

$$l(a, b) \triangleq -a \log(b) - (1 - a) \log(1 - b). \quad (32)$$

In this case, using the FL strategy in (15), we obtain the same closed form solution as in (19) (or (26)) by using the ℓ_2 (or ℓ_1) regularization, with the only difference being that each function g_i in (4) now depends on the cross-entropy loss in (32). The PL case, instead, requires some additional care. In particular, although the NN output is bounded, the same is not true for its linear approximation (16). Simply substituting (16) in (32) might result in undefined values, since the argument of the logarithm must be positive. To tackle this issue, let us notice that in this case the NN model can be written

as:

$$f(\mathbf{w}; \mathbf{x}) = \sigma(f_L(\mathbf{w}; \mathbf{x})) , \quad (33)$$

where $\sigma(\cdot)$ is a squashing function (without loss of generality, we assume it to be a sigmoid), and f_L is the NN output up to (but not including) the activation function of the output neuron. The sigmoid $\sigma(z)$ is non-convex, but its internal composition with the cross-entropy loss in (32) is convex, see, e.g., (Boyd and Vandenberghe, 2004). Exploiting such hidden convexity, we can write the surrogate problem in (13), while satisfying the conditions F1-F3 in Proposition 1, as follows:

$$\boxed{\begin{aligned} \widetilde{\mathbf{w}}_i[n] = \arg \min_{\mathbf{w}_i} & \left\{ l \left(d_{i,m}, \sigma \left(\widetilde{f}_L(\mathbf{w}_i; \mathbf{w}_i[n], \mathbf{x}_{i,m}) \right) \right) \right. \\ & \left. + \boldsymbol{\pi}_i[n]^T \mathbf{w}_i + r(\mathbf{w}_i) + \frac{\tau}{2} \|\mathbf{w}_i - \mathbf{w}_i[n]\|_2^2 \right\} , \end{aligned}} \quad (34)$$

where $\widetilde{f}_L(\cdot)$ is the first-order linearization of f_L defined as in (16). Also in this case we cannot make any further simplifications although, once again, the strong convexity of the problem makes it relatively easy to be solved (roughly equivalent to a traditional logistic regression).

5. Parallelizing the Local Optimization

In this section, we explore how each agent can parallelize the local optimization in (13), when having access to C separate computing machines (e.g. cores, or computers in a cloud). As we stated in the introduction, this effectively gives rise to algorithms that are both distributed (across agents) and parallel (inside each agent) in nature. To this end, suppose that the weight vector \mathbf{w}_i is partitioned in C non-overlapping blocks $\mathbf{w}_{i,1}, \dots, \mathbf{w}_{i,C}$, so that $\mathbf{w}_i = \bigcup_{c=1}^C \mathbf{w}_{i,c}$. Note that we use a similar notation as in Section 4.3 to identify a single group, i.e., using an additional subscript under the variable. For convenience, we also define $\mathbf{w}_{i,-c}[n] \triangleq (\mathbf{w}_{i,p}[n])_{1=p \neq c}^C$ as the tuple of all blocks excepts the c -th one, and similarly for all other variables. Additionally, we assume that the regularization term r is block separable, i.e., $r(\mathbf{w}_i) = \sum_{c=1}^C r_{i,c}(\mathbf{w}_{i,c})$ for some $r_{i,c}$. This is true for the ℓ_2 and ℓ_1 norms, and it holds true also for the group sparse norm in (30) if we choose the groups in a consistent way. Then, the key idea is to decompose (13) on a

per-core-basis, and solve a sequence of (strongly) convex low-complexity sub-problems, whereby all processors of agent i update their blocks in parallel. To this aim, we build a surrogate fashion \tilde{g}_i that additively decomposes over the different cores, i.e.:

$$\tilde{g}_i(\mathbf{w}_i; \mathbf{w}_i[n]) = \sum_{c=1}^C \tilde{g}_{i,c}(\mathbf{w}_{i,c}; \mathbf{w}_{i,-c}[n]), \quad (35)$$

where each $\tilde{g}_{i,c}(\cdot; \mathbf{w}_{i,-c}[n])$ is any surrogate function satisfying conditions F1-F3 on the variable $\mathbf{w}_{i,c}$. It is easy to check that the surrogate \tilde{g}_i in (35) satisfies F1-F3 on the variable \mathbf{w}_i . Given (35), each core c can then minimize its corresponding term independently of the others, and their solutions can be aggregated to form the final solution vector. In the case of the FL strategy, parallelization is not particularly effective. In fact, the final solution is given by simple aggregation of vectors as in (19), whose computation has linear complexity with respect to the size of \mathbf{w}_i , eventually with a point-wise application of the thresholding operator in (28). However, the (linear) cost of solving the surrogate problems at each core is easily overshadowed by the need of computing gradients via a backpropagation step. On the other side, parallelization can largely reduce computational complexity when using the PL strategy. To give an example of application of the proposed methodology, in the sequel we illustrate how to parallelize the local optimization in the case of a ridge regression cost as in Sec. 4.1. In particular, let us consider the surrogate function in (21). To obtain the surrogate function associated to each core c , we fix in (21) all the variables $\mathbf{w}_{i,-c}[n]$, such that the resulting function depends only on $\mathbf{w}_{i,c}$. The surrogate associated to core c is then given by:

$$\begin{aligned} \tilde{U}_{i,c}(\mathbf{w}_{i,c}; \mathbf{w}_{i,-c}[n]) = & \mathbf{w}_{i,c}^T \left(\mathbf{A}_{i,c,c}[n] + \frac{\lambda}{2} \mathbf{I} \right) \mathbf{w}_{i,c} - \\ & 2 \left(\mathbf{b}_{i,c}[n] - \mathbf{A}_{i,c,-c}[n] \mathbf{w}_{i,-c}[n] \right)^T \mathbf{w}_{i,c}, \end{aligned} \quad (36)$$

where $\mathbf{A}_{i,c,c}[n]$ is the block (rows and columns) of the matrix $\mathbf{A}_i[n]$ in (22) corresponding to the c -th partition, whereas $\mathbf{A}_{i,c,-c}[n]$ takes the rows corresponding to the c -th partition and all the columns not associated to c . The

minimum of (36) is:

$$\tilde{\mathbf{w}}_{i,c}[n] = \left(\mathbf{A}_{i,c,c}[n] + \frac{\lambda}{2} \mathbf{I} \right)^{-1} (\mathbf{b}_{i,c}[n] - \mathbf{A}_{i,c,-c}[n] \mathbf{w}_{i,-c}[n]) , \quad (37)$$

$c = 1, \dots, C$, and the overall solution is given by $\tilde{\mathbf{w}}_i[n] = (\tilde{\mathbf{w}}_{i,c}[n])_{c=1}^C$. As we can see from (37), the effect of the parallelization is evident: At each iteration n , each core has to invert a matrix having (approximately) size $\frac{1}{C}$ of the original one in (24), thus remarkably reducing the overall computational burden. Similar arguments can be used also to parallelize the formulations in (29), (31), and (34).

6. Experimental Validation

In this section, we assess the performance of the proposed method via numerical simulations. We begin by analyzing the test error of the solutions obtained by the algorithms for some representative regression and classification datasets in Sections 6.2 and 6.3, respectively. Then, we consider the convergence behaviors of the proposed framework, comparing it to centralized and distributed counterparts, in Section 6.4. Finally, in Section 6.5, we describe the speed-up achieved thanks to the parallelization strategy outlined before. Python code to repeat the experiments is available under open-source license on the web.⁴ The code is built upon the Theano (Bergstra et al., 2010) and Lasagne⁵ libraries.

6.1. Experimental setup

In all experiments, the original dataset is normalized so that both inputs and outputs lies in the $[0, 1]$ range. Then, the dataset is partitioned as follows. First, 20% of the dataset is kept separate to test the algorithms. The remaining 80% is partitioned evenly among a randomly generated network of 10 agents. For simplicity, we consider networks with time-invariant, symmetric connectivity, such that every pair of agents have a 20% probability of being connected, with the only requirement that the overall network is connected. An example of such connectivity is shown in Fig. 1.

⁴<https://bitbucket.org/ispamm/parallel-and-distributed-neural-networks>

⁵<https://github.com/Lasagne/Lasagne>




Figure 1: Example of communication network with 10 agents (represented by red dots), possessing a sparse, time-invariant, symmetric connectivity.

We have selected the weight coefficients in (2) using the *Metropolis-Hastings* strategy (Lopes and Sayed, 2008):

$$C_{ij} = \begin{cases} \frac{1}{\max\{d_i, d_j\}+1} & i \neq j, j \in \mathcal{N}_i \\ 1 - \sum_{j \in \mathcal{N}_i} \frac{1}{\max\{d_i, d_j\}+1} & i = j \\ 0 & i \neq j, \{i, j\} \notin \mathcal{N}_k \end{cases} \quad (38)$$

where d_i is the degree of node i . It is easy to check that this choice of the weight matrix satisfies the convergence conditions of the framework. Missing data is handled by removing the corresponding example. All experiments are repeated 25 times by varying the data partitioning and the NN initialization.

Regarding the NN structure, we use hyperbolic tangent nonlinearities in all neurons, except for classification problems, where we use a sigmoid nonlinearity in the output neuron. The weights of the NN are initialized independently at every agent using the normalized strategy described in Glorot and Bengio (2010). All algorithms run for a maximum of 1000 epochs. In all the figures illustrating the results of the distributed strategies, whenever not explicitly stated, we consider the evolution of the average weight vector $\bar{\mathbf{w}}[n]$ as defined in Theorem 1.

6.2. Results with regression datasets

We start considering three representative regression datasets, whose characteristics are summarized in Table 1. Boston (also known as the Housing dataset) is the task of predicting the median value of a house based on a set

Dataset	Features	Samples	NN Topology	λ	Source
Boston	13	506	10	10^{-1}	UCI
Kin8nm	7	8192	8/5	10^{-2}	DELVE
Wine	10	4898	12	10^{-2}	UCI

Table 1: Schematic description of the datasets used for regression. For the NN topology, x/y denotes a NN with two hidden layers of dimensions x and y respectively.

of features describing it (Quinlan, 1993). Kin8nm is a member of the kinematics family of datasets⁶, having high non-linearity and a medium amount of additive noise. Finally, Wine concerns predicting the subjective quality of a (white) wine based on a wide set of chemical features (Cortez et al., 2009). The fourth and fifth columns in Table 1 describe the parameters of the NN in terms of hidden neurons and regularization coefficients. These parameters are chosen based on an analysis of previous literature in order to obtain state-of-the-art results. However, we underline that our aim is to compare different solvers for the same NN optimization problem, and for this reason only relative differences in accuracy are of concern.

In particular, we compare the results of our algorithms with respect to five state-of-the-art *centralized* solvers, in terms of mean-squared error (MSE) over the test data, when solving the global optimization problem with the ridge regression cost in (18). Note that these solvers would not be available in a distributed scenario, and are only used for comparison purposes as optimal benchmarks. Specifically, we consider the following algorithms:

Gradient descent (GD) : this is a simple first-order steepest descent procedure with fixed step-size.

AdaGrad (Duchi et al., 2011) : differently from GD, this algorithm employs different adaptive step-sizes per weight, which evolve according to the relative values of the gradients’ updates.

RMSProp : equivalent to an AdaDelta variant (Zeiler, 2012), it also considers adaptive independent step-sizes; however they are adapted based

⁶<http://www.cs.toronto.edu/~delve/data/kin/desc.html>

on shorter time windows in order to avoid exponentially decreasing schedules.

Conjugate gradient (CG) : this is the Polak-Ribiere variant of the non-linear conjugate gradient algorithm (Nocedal and Wright, 2006), implemented in the SciPy library.⁷

L-BFGS : a low-memory version of the second-order BFGS algorithm (Byrd et al., 1995), keeping track of an approximation to the full Hessian matrix, also implemented in the SciPy library.

In addition, we consider the behavior of a centralized implementation of the PL strategy in (24), denoted as PL-SCA, resulting in a novel centralized algorithm. In particular, assuming all data is available on a single location, we can consider a centralized equivalent of (22) and (23) as:

$$\mathbf{A}[n] = \sum_{i=1}^I \mathbf{A}_i[n], \quad (39)$$

$$\mathbf{b}[n] = \sum_{i=1}^I \sum_{m \in S_i} \mathbf{J}_{i,m}[n] r_{i,m}[n]. \quad (40)$$

Following similar arguments as in Sections 3 and 4.1, PL-SCA is defined by the iterative application of the following recursion:

$$\tilde{\mathbf{w}}[n] = \left(\mathbf{A}[n] + \frac{\lambda}{2} \mathbf{I} \right)^{-1} \mathbf{b}[n], \quad (41)$$

$$\mathbf{w}[n+1] = \mathbf{w}[n] + \alpha[n] (\tilde{\mathbf{w}}[n] - \mathbf{w}[n]). \quad (42)$$

The centralized implementation of the FL strategy is almost equivalent to GD, so we do not consider it separately. For the distributed algorithms, we consider both the PL strategy in (24), denoted as PL-NEXT, and the FL strategy in (19), denoted as FL-NEXT. For PL-SCA, PL-NEXT and FL-NEXT we use the quadratically decreasing step-size sequence defined in (14). To have a fair comparison, the parameters of the step-size sequence in (14) were tuned at hand in order to select the fastest convergence behavior for all algorithms.

⁷<https://scipy.org/>

Dataset	Centralized					Distributed	
	GD	AdaGrad	RMSProp	CG	L-BFGS	PL-SCA	FL-NEXT PL-NEXT
Boston	0.010 ± 0.001	0.009 ± 0.001	0.009 ± 0.001	0.007 ± 0.001	0.007 ± 0.001	0.007 ± 0.001	0.010 ± 0.001 0.007 ± 0.001
Kin8nm	0.019 ± 0	0.018 ± 0	0.015 ± 0.001	0.011 ± 0.001	0.009 ± 0.001	0.009 ± 0	0.019 ± 0 0.009 ± 0
Wine	0.018 ± 0.001	0.016 ± 0.001	0.017 ± 0.001	0.038 ± 0.019	0.014 ± 0	0.014 ± 0.001	0.017 ± 0.001 0.014 ± 0.001

Table 2: Regression results (in terms of mean-squared error on the test set) of different models. Columns 2-7 are centralized algorithms reported for comparison. Columns 8-9 are the two implementations of the proposed distributed framework. Best results for both groups are highlighted in bold. See the text for a description of the acronyms.

The results on this set of experiments are provided in Table 2, both in terms of mean and standard deviation. Several conclusions can be drawn from the table. For the centralized algorithms, L-BFGS, being a second-order algorithm, is able to obtain the best accuracies, and it is matched only by CG in the Boston case (in the next section we will show some plots of the convergence behavior of the different algorithms). Interestingly, PL-SCA is able to match L-BFGS in all cases, complementing our previous observation that the matrix in (39) acts as an approximation of the Hessian matrix. For the distributed algorithms, we see similar distinctions between FL-NEXT and PL-NEXT. Specifically, PL-NEXT has comparable accuracies with respect to L-BFGS, while FL-NEXT obtains errors comparable to GD and AdaGrad. Clearly, the improved convergence comes at the cost of a higher computational burden, in line with the equivalent difference in the centralized case. Summarizing, we see that FL-NEXT and PL-NEXT represent viable algorithms for distributed scenarios, providing a relative trade-off with respect to convergence and computational requirements, and matching the respective centralized implementations that are not viable in the distributed setting treated in this papers. Importantly, this is also achieved with a minimal (or non-existent) increase in term of variance.

6.3. Results with classification datasets

In this section, we analyze the performance of the distributed algorithms when applied to two classification problems, whose characteristics are briefly summarized in Table 3. Wisconsin is a medical classification task, aimed at separating cancerous cells from non-cancerous ones from several features describing the cell nucleus.⁸ The Cardiotocography (CGT) is another clinical problem, where we wish to infer suspect/pathological fetuses from several biometric signals.⁹ In this case, we solve the global optimization problem with the cross-entropy loss in (32) and a squared regularization term. We analyze the behavior of both the FL strategy and the PL strategy when compared to the state-of-the-art solvers described in the previous section. For PL-NEXT, the local surrogate problem in (34) is solved with AdaGrad, run for a maximum of 50 iterations (with an initial step-size of 0.1), or until the

⁸[https://archive.ics.uci.edu/ml/datasets/Breast+Cancer+Wisconsin+\(Diagnostic\)](https://archive.ics.uci.edu/ml/datasets/Breast+Cancer+Wisconsin+(Diagnostic))

⁹<https://archive.ics.uci.edu/ml/datasets/Cardiotocography>

Dataset	Features	Samples	NN Topology	λ	Source
Wisconsin	9	689	10	$10^{-0.5}$	UCI
CTG	28	2126	15/8	10	UCI

Table 3: Schematic description of the datasets used for classification. See Table 3 and the text for details on the NN topology.

gradient norm is below a fixed threshold of 10^{-6} . For the local optimization at each agent, we perform a ‘warm start’ from the current estimate $\mathbf{w}_i[n]$.

The overall results are given in Table 4 in terms of misclassification rate. We see that, in this case, first-order algorithms are generally competitive, with the GD solver obtaining the best accuracy among the centralized solvers for the Wisconsin dataset, and CG/L-BFGS obtaining a slightly better result in the CTG case. Nevertheless, the distributed strategies are again able to obtain state-of-the-art results, with PL-NEXT consistently obtaining the lowest misclassification rate, and FL-NEXT ranging close to AdaGrad and RMSProp.

6.4. Analysis of convergence

In a distributed setting, the final accuracy is not the only parameter of interest. We are also concerned on how fast this accuracy is obtained, because the convergence speed has a direct impact on the communication burden over the network of agents. As we mentioned in the introduction, in the case of general non-differentiable regularizers r , there is no ready-to-use alternative for comparing our proposed algorithms. However, in the specific case where the regularization function r satisfies assumption B1, we can easily adapt the framework introduced in Bianchi and Jakubowicz (2013), resulting in a simple method that we denote as the distributed gradient (DistGrad) algorithm. Similarly to the NEXT framework, DistGrad alternates between a local optimization phase and a communication phase. In the optimization phase, each agent iteratively updates its own estimate according to a local gradient descent step as follows:

$$\mathbf{z}_i[n] = \mathbf{w}_i[n] - \eta[n] \left(\nabla \mathbf{g}_i[n] + \frac{1}{I} \nabla r(\mathbf{w}_i[n]) \right), \quad (43)$$

Dataset	Centralized				Distributed	
	GD	AdaGrad	RMSprop	CG	L-BFGS	FL-NEXT PL-NEXT
Wisconsin	0.025 ± 0.009	0.027 ± 0.006	0.027 ± 0.005	0.028 ± 0.006	0.028 ± 0.006	0.027 ± 0.007 0.025 ± 0.009
CTG	0.084 ± 0.010	0.082 ± 0.010	0.087 ± 0.014	0.083 ± 0.011	0.083 ± 0.010	0.087 ± 0.007 0.084 ± 0.009

Table 4: Classification results (in terms of misclassification rate on the test set) of different models. Columns 2-6 are centralized algorithms reported for comparison. Columns 7-8 are the two implementations of the proposed distributed framework. Best results for both groups are highlighted in bold. See the text for a description of the acronyms.

where $\eta[n]$ is the step-size sequence. In the communication phase, the local estimates $\mathbf{z}_i[n]$ are combined similarly to (8). DistGrad can be seen as a simplified version of the FL strategy, where we do not consider the dynamic consensus step (i.e., Step 3 of NEXT). For fairness of comparison, we use the step-size rule in (14), and the max-degree strategy for selecting the combination coefficients in (2).


In Figs. 2a-2b we plot the evolution of the global cost function in (3) for FL-NEXT, PL-NEXT, DistGrad and a few representative centralized solvers for two different datasets. For improved readability, the behavior of centralized solvers is depicted using dashed lines, while the distributed algorithms are shown with solid lines. We see that the results are similar to what we have already discussed previously for the final test error: PL-NEXT is able to track consistently the convergence rate of L-BFGS, while FL-NEXT achieves results comparable to (centralized) first-order procedures. Differently, the DistGrad algorithm is slower and, for a given number of epochs, has a very large gap compared to other methods.

Another performance metric of interest is the transient behavior of the test error in terms of the amount of scalar values that are exchanged among agents in the network. We plot this metric for the three distributed algorithms in Figs. 2c-2d, where the y -axis is shown with a logarithmic scale. We notice that DistGrad requires exactly half as many scalars to be exchanged at every iteration (since it does not rely on the dynamic consensus to track the average gradient). Nevertheless, from Figs. 2c-2d, we can see that both PL-NEXT and FL-NEXT can reach better errors with respect to DistGrad for any given amount of scalars exchanged, showing their better efficiency in terms of overall communication burden. PL-NEXT is particularly well performing, with only a very small amount of communication required for achieving an error close to the optimal one.


A final metric of interest is the average disagreement among the agents, which is computed as:

$$D[n] \triangleq \frac{1}{I} \|\mathbf{w}_i[n] - \bar{\mathbf{w}}[n]\|_\infty . \quad (44)$$


We plot the behavior of (44) for PL-NEXT and FL-NEXT in Figs. 2e-2f, where we can see that both algorithms rapidly tend to reach a consensus among the different agents in the network.




(a) Cost function (Boston)




(b) Cost function (Kin8nm)




(c) Test error (Boston)



(d) Test error (Kin8nm)



(e) Disagreement (Boston)



(f) Disagreement (Kin8nm)

Figure 2: (a-b) Cost function value (per epoch). (c-d) Test error (per scalars exchanged). (e-f) Evolution of the disagreement (per scalars exchanged). Graphs on the left column are for the Boston dataset, graphs on the right column are for the Kin8nm dataset. For readability, centralized algorithms are represented with dashed lines, while distributed algorithms are represented with solid lines with specific markers.

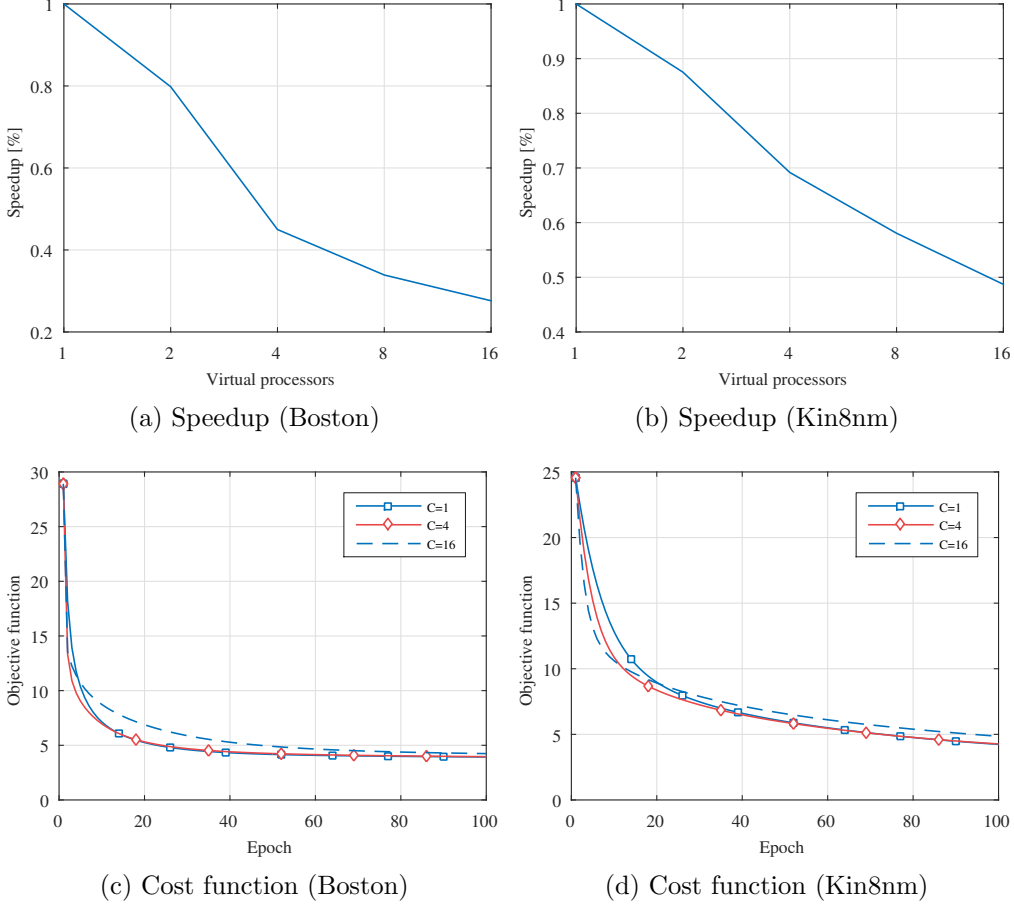


Figure 3: (a-b) Relative speedup (per number of local processors). (c-d) Cost function value (per epoch). Graphs on the left column are for the Boston dataset, graphs on the right column are for the Kin8nm dataset.

6.5. Exploiting parallelization

Finally, we investigate the speed-up obtained by parallelizing the local optimization at each agent. We consider again the Boston and Kin8nm datasets, but we vary the number of (local) processors available at every agent in the range 2^j , with $j = 0, 1, \dots, 4$. The relative speedup with respect to the baseline $C = 1$ is shown in Figs. 3a-3b. We see that the speedup is roughly linear with respect to the amount of available processors, so that in the case $C = 16$ we only need $\approx \frac{1}{3}$ of the time for Boston, and $\approx \frac{1}{2}$ for Kin8nm. Additionally, in Figs. 3c-3d, we can visualize the evolution of the

overall cost function for $C = 1$, $C = 4$ and $C = 16$. From the figures, we can notice that the improvement in training time is obtained with only a limited effect on the convergence behavior, where in the worst case we obtain only a very small (or null) decrease.

7. Conclusions and Future Works

In this paper, we have investigated the problem of training a NN model in a distributed scenario, where multiple agents have a limited knowledge of the training dataset. We have proposed a provably convergent procedure to this aim, which builds exclusively on local optimization steps and one-hop communication steps. The method can be customized to several typical error functions and regularization terms. We have also described an immediate way to parallelize the local optimization phase across multiple processors/machines, available at each agent, with a limited impact on the convergence behavior.

One immediate extension of the framework presented here is to handle non-convex regularization terms, which are generally considered too challenging in practice. One example is the sample variance penalization ([Maurer and Pontil, 2009](#)), which is defined in terms of the NN output. Additional extensions can consider the presence of non-differentiable points in the NN model (e.g. by using ReLu activation functions), or stochastic updates of the surrogate functions. Finally, we are interested in testing our framework on real-world applications such as, e.g., multimedia classification and chaotic prediction tasks.

References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., et al., 2016. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. arXiv preprint arXiv:1603.04467.
- Beck, A., Teboulle, M., 2009. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences* 2 (1), 183–202.
- Bengio, Y., 2012. Practical recommendations for gradient-based training of deep architectures. In: *Neural Networks: Tricks of the Trade*. Springer, pp. 437–478.

- Bergstra, J., Breuleux, O., Bastien, F., Lamblin, P., Pascanu, R., Desjardins, G., Turian, J., Warde-Farley, D., Bengio, Y., 2010. Theano: A cpu and gpu math compiler in python. In: Proceedings of the 9th Python in Science Conference. pp. 1–7.
- Bianchi, P., Jakubowicz, J., 2013. Convergence of a multi-agent projected stochastic gradient algorithm for non-convex optimization. *IEEE Transactions on Automatic Control* 58 (2), 391–405.
- Bishop, C. M., 2006. Pattern recognition and machine learning. Springer International.
- Blackwell, W. J., 2012. Neural network jacobian analysis for high-resolution profiling of the atmosphere. *EURASIP Journal on Advances in Signal Processing* 2012 (1), 1.
- Boric-Lubeke, O., Lubecke, V. M., 2002. Wireless house calls: using communications technology for health care and monitoring. *IEEE Microwave Magazine* 3 (3), 43–48.
- Boyd, S., Vandenberghe, L., 2004. Convex optimization. Cambridge university press.
- Byrd, R. H., Lu, P., Nocedal, J., Zhu, C., 1995. A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing* 16 (5), 1190–1208.
- Cevher, V., Becker, S., Schmidt, M., 2014. Convex optimization for big data: Scalable, randomized, and parallel algorithms for big data analytics. *IEEE Signal Processing Magazine* 31 (5), 32–43.
- Cortez, P., Cerdeira, A., Almeida, F., Matos, T., Reis, J., 2009. Modeling wine preferences by data mining from physicochemical properties. *Decision Support Systems* 47 (4), 547–553.
- Dean, J., Corrado, G., Monga, R., Chen, K., Devin, M., Mao, M., Senior, A., Tucker, P., Yang, K., Le, Q. V., et al., 2012. Large scale distributed deep networks. In: Advances in neural information processing systems. pp. 1223–1231.

- Di Lorenzo, P., Sayed, A. H., 2013. Sparse distributed learning based on diffusion adaptation. *IEEE Transactions on Signal Processing* 61 (6), 1419–1433.
- Di Lorenzo, P., Scardapane, S., 2016. Parallel and distributed training of neural networks via successive convex approximation. In: *2016 IEEE International Workshop on Machine Learning for Signal Processing (MLSP)*. IEEE, pp. 1–6.
- Di Lorenzo, P., Scutari, G., 2016. Next: In-network nonconvex optimization. *IEEE Transactions on Signal and Information Processing over Networks* 2 (2), 120–136.
- Duchi, J., Hazan, E., Singer, Y., 2011. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research* 12 (Jul), 2121–2159.
- Facchinei, F., Scutari, G., Sagratella, S., 2015. Parallel selective algorithms for nonconvex big data optimization. *IEEE Transactions on Signal Processing* 63 (7), 1874–1889.
- Forero, P. A., Cano, A., Giannakis, G. B., 2010. Consensus-based distributed support vector machines. *Journal of Machine Learning Research* 11 (May), 1663–1707.
- Gao, W., Chen, J., Richard, C., Huang, J., 2015. Diffusion adaptation over networks with kernel least-mean-square. In: *Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP), 2015 IEEE 6th International Workshop on*. IEEE, pp. 217–220.
- Georgopoulos, L., Hasler, M., 2014. Distributed machine learning in networks by consensus. *Neurocomputing* 124, 2–12.
- Glorot, X., Bengio, Y., 2010. Understanding the difficulty of training deep feedforward neural networks. In: *AISTATS*. Vol. 9. pp. 249–256.
- Glorot, X., Bordes, A., Bengio, Y., 2011. Deep Sparse Rectifier Neural Networks. In: *Proc. 14th International Conference on Artificial Intelligence and Statistics*. pp. 315–323.

- Goodfellow, I. J., Warde-farley, D., Mirza, M., Courville, A., Bengio, Y., 2013. Maxout networks. In: Proc. 30th International Conference on Machine Learning. pp. 1319–1327.
- Haykin, S., 2009. Neural networks and learning machines, 3rd Edition. Pearson.
- Huang, S., Li, C., 2015. Distributed extreme learning machine for nonlinear learning over network. *Entropy* 17 (2), 818–840.
- Lazarevic, A., Obradovic, Z., 2002. Boosting algorithms for parallel and distributed learning. *Distributed and Parallel Databases* 11 (2), 203–229.
- LeCun, Y., Bengio, Y., Hinton, G., 2015. Deep learning. *Nature* 521 (7553), 436–444.
- Lopes, C. G., Sayed, A. H., 2008. Diffusion least-mean squares over adaptive networks: Formulation and performance analysis. *IEEE Transactions on Signal Processing* 56 (7), 3122–3136.
- Lu, Y., Roychowdhury, V., Vandenberghe, L., 2008. Distributed parallel support vector machines in strongly connected networks. *IEEE Transactions on Neural Networks* 19 (7), 1167–1178.
- Mateos, G., Bazerque, J. A., Giannakis, G. B., 2010. Distributed sparse linear regression. *IEEE Transactions on Signal Processing* 58 (10), 5262–5276.
- Maurer, A., Pontil, M., 2009. Empirical bernstein bounds and sample variance penalization. *arXiv preprint arXiv:0907.3740*.
- Moody, J., Hanson, S., Krogh, A., Hertz, J. A., 1995. A simple weight decay can improve generalization. *Advances in neural information processing systems* 4, 950–957.
- Navia-Vázquez, A., Gutierrez-Gonzalez, D., Parrado-Hernández, E., Navarro-Abellan, J., 2006. Distributed support vector machines. *IEEE Transactions on Neural Networks* 17 (4), 1091–1097.
- Nocedal, J., Wright, S., 2006. Numerical optimization. Springer Science & Business Media.

- Ochs, P., Dosovitskiy, A., Brox, T., Pock, T., 2015. On iteratively reweighted algorithms for nonsmooth nonconvex optimization in computer vision. *SIAM Journal on Imaging Sciences* 8 (1), 331–372.
- Perez-Cruz, F., Kulkarni, S. R., 2010. Robust and low complexity distributed kernel least squares learning in sensor networks. *IEEE Signal Processing Letters* 17 (4), 355–358.
- Pottie, G. J., Kaiser, W. J., 2000. Wireless integrated network sensors. *Communications of the ACM* 43 (5), 51–58.
- Predd, J., Kulkarni, S., Poor, H., 2006. Distributed learning in wireless sensor networks. *IEEE Signal Processing Magazine* 23 (4), 56–69.
- Predd, J. B., Kulkarni, S. R., Poor, H. V., 2009. A collaborative training algorithm for distributed learning. *IEEE Transactions on Information Theory* 55 (4), 1856–1871.
- Quinlan, J. R., 1993. Combining instance-based and model-based learning. In: *Proceedings of the Tenth International Conference on Machine Learning*. pp. 236–243.
- Sak, H., Vinyals, O., Heigold, G., Senior, A., McDermott, E., Monga, R., Mao, M., 2014. Sequence discriminative distributed training of long short-term memory recurrent neural networks. In: *Interspeech 2014*.
- Samet, S., Miri, A., 2012. Privacy-preserving back-propagation and extreme learning machine algorithms. *Data & Knowledge Engineering* 79, 40–61.
- Sayed, A. H., 2014. Adaptive networks. *Proceedings of the IEEE* 102 (4), 460–497.
- Sayed, A. H., et al., 2014. Adaptation, learning, and optimization over networks. *Foundations and Trends® in Machine Learning* 7 (4-5), 311–801.
- Scardapane, S., Comminiello, D., Hussain, A., Uncini, A., 2016a. Group sparse regularization for deep neural networks. *arXiv preprint arXiv:1607.00485*.
- Scardapane, S., Fierimonte, R., Di Lorenzo, P., Panella, M., Uncini, A., 2016b. Distributed semi-supervised support vector machines. *Neural Networks* 80, 43–52.

- Scardapane, S., Wang, D., Panella, M., 2016c. A decentralized training algorithm for echo state networks in distributed big data applications. *Neural Networks* 78, 65–74.
- Scardapane, S., Wang, D., Panella, M., Uncini, A., 2015. Distributed learning for random vector functional-link networks. *Information Sciences* 301, 271–284.
- Schmidhuber, J., 2015. Deep learning in neural networks: An overview. *Neural Networks* 61, 85–117.
- Schmidt, M., 2010. Graphical model structure learning with l1-regularization. Ph.D. thesis, The University of British Columbia (Vancouver).
- Tibshirani, R., 1996. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 267–288.
- Vieira-Marques, P. M., Robles, S., Cucurull, J., Navarro, G., et al., 2006. Secure integration of distributed medical data using mobile agents. *IEEE Intelligent Systems* (6), 47–54.
- Xiao, L., Boyd, S., 2004. Fast linear iterations for distributed averaging. *Systems & Control Letters* 53 (1), 65–78.
- Xiao, L., Boyd, S., Kim, S.-J., 2007. Distributed average consensus with least-mean-square deviation. *Journal of Parallel and Distributed Computing* 67 (1), 33–46.
- Zeiler, M. D., 2012. Adadelta: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701*.
- Zhang, Y., Zhong, S., 2013. A privacy-preserving algorithm for distributed training of neural network ensembles. *Neural Computing and Applications* 22 (1), 269–282.
- Zhu, M., Martínez, S., 2010. Discrete-time Dynamic Average Consensus. *Automatica* 46 (2), 322–329.